REMARKS/ARGUMENTS

Claims 1-4, 6-8, 10-14, 16-18, 20-24, 26-28, and 30 are pending in the application. Claims 1, 7, 11, 17, 21, and 27 have been amended merely to overcome 35 U.S.C. 112 and 101 rejections. Claims 5, 9, 15, 19, 25, and 29 have been cancelled. Reconsideration is respectfully requested. Applicants submit that the pending claims 1-4, 6-8, 10-14, 16-18, 20-24, 26-28, and 30 are patentable over the art of record and allowance is respectfully requested of claims 1-4, 6-8, 10-14, 16-18, 20-24, 26-28, and 30.

Applicants would like to thank Examiners Mahmoudi and Radtke for holding a telephone interview with their representative, Janaki K. Davda, on December 5, 2006 at 3:00 p.m. (EST). Proposed claim amendments to claim 1 were discussed. Examiner Mahmoudi indicated that a new search would be required.

Claims 1-4, 6-8, 10-14, 16-18, 20-24, 26-28, and 30 are rejected under 35 U.S.C. 112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which applicant regards as the invention. Applicants respectfully traverse, but, in order to expedite prosecution, Applicants have amended certain claims.

As to claims 1-4, 6-8, 10-14, 16-18, 20-24, 26-28, and 30, the Examiner indicates that the use of the term "input duplicates" is considered to lack antecedent basis. Applicants respectfully traverse. The first time the term "input duplicates" is introduced, there is no antecedent basis required (as there is no antecedent use of the term), and subsequent uses of the term "input duplicates" are preceded by "the" to provide the proper antecedent basis. Additionally, Applicants are clarifying the term "input duplicates" as duplicate input rows having a same primary key value, and wherein, of the input rows having the same primary key value, one is an original row (e.g., Specification, page 13, paragraph 42). Thus, the Examiner's assumption that these are "entries that were classified as update rows" is incorrect.

As to claims 1-4, 6-8, 10-14, 16-18, 20-24, 26-28, and 30, the Examiner indicates that the term "unique index entries" is considered to lack antecedent basis. The Examiner submits that no suitable interpretation of this phrase can be made because only keys corresponding to input duplicates are stored in the second structure and so it would be impossible for any unique keys to

be stored there. Applicants respectfully traverse, but, in order to expedite prosecution, Applicants have removed the term "unique".

Claims 1-4, 6-8, 10-14, 16-18, 20-24, 26-28, and 30 are rejected under 35 U.S.C. 101 because the disclosed invention is inoperative and therefore lacks utility. Applicants respectfully traverse.

The Examiner is interpreting an insert row as a non-duplicate row and an update row as a duplicate row. Applicants submit that this is a misinterpretation. Applicants have amended the claims to clarify that an insert row is a row to be added to the first structure and an update row is an update to an existing row in the first structure (e.g., Specification, page 1, paragraph 3). Also, Applicants have amended the claims to clarify that input duplicates are duplicate input rows having a same primary key value, and wherein, of the input rows having the same primary key value, one is an original row (e.g., Specification, page 13, paragraph 42).

Also, the Examiner indicates that "most databases strictly enforce unique values for primary keys and this step would append duplicate entries with duplicate keys", which would render the invention inoperable. Applicants respectfully submit that Applicants are allowed to invent new techniques, including ones that append input duplicates to an output table. The Examiner's use of the term "most" allows for some tables that do not enforce unique values for primary keys. Also, the Examiner's argument goes to the novelty of Applicants' claimed invention.

The Examiner submits that re-applying input duplicates is "redundant and produces no net effect on the output table". Applicants respectfully traverse. Applicants have amended the claims to indicate that automatically re-applying the input duplicates to the first structure by removing the input duplicates from the first structure and applying the input duplicates to matching original rows in the first structure to update the matching original rows (e.g., Specification, page 13, paragraph 43; page 21, paragraph 72). Thus, the re-application updates the original rows. Note that an earlier amended claim element clarifies that input duplicates are duplicate input rows having a same primary key value, and wherein, of the input rows having the same primary key value, one is an original row.

The Examiner submits that merging index entries is illustrated in figure 6A of Applicants' Specification and that it is redundant. Applicants respectfully traverse. Figure 6A of Applicants'

Specification describes processing with reference to the output table (e.g., Specification, pages 14-15, paragraphs 47-48). Also, the claim specifically states "merging index entries from the index entries stored in the second structure to a primary key index". That is, the index entries stored in the second structure are not merged into the output table, but are merged into the primary key index, which is another structure. For example, in figure 6B of Applicants' Specification, index A 690 is a primary key index (e.g., Specification, page 15, paragraph 49). Thus, the merging of index entries is not redundant.

Claims 1-4, 6-8, 10-14, 16-18, 20-24, 26-28, and 30 are rejected under 35 U.S.C. 103(a) as being unpatentable over Thusoo et al. (U.S. Patent No. 7,016,903) in view of Oracle ("MERGE Statement"). Applicants respectfully traverse.

Claims 1, 11, and 21 describe receiving multiple input rows to be loaded into a first structure, wherein the first structure is an output table (e.g., Specification, page 12, paragraphs 40-41). Each input row of the multiple input rows is processed to classify each input row as one of an insert row and an update row, wherein input duplicates are appended to the first structure and index entries for the input duplicates are stored in a second structure (e.g., Specification, page 12, paragraph 41; FIG. 5A, block 512). An insert row is a row to be added to the first structure, and an update row is an update to an existing row in the first structure. Input duplicates are duplicate input rows having a same primary key value, and, of the input rows having the same primary key value, one is an original row. After the multiple input rows have been processed, the input duplicates are automatically re-applied to the first structure by removing the input duplicates from the first structure and applying the input duplicates to matching original rows in the first structure to update the matching original rows (e.g., Specification, page 13, paragraph 43; FIG. 5B, blocks 522, 524) and index entries from the index entries stored in the second structure are merged to a primary key index (e.g., Specification, page 13, paragraph 43; FIG. 5, block 518).

On the other hand, the Thusoo patent describes that to process the source table 305 and integrate it with destination table 310, the data warehouse 300 may completely join the two tables together. . . and the result of the outer join step is preferably a result set 340 (Col. 8, lines 1-11). Once result set 340 has been compiled, each row in result set 340 may be processed such

that the row is either inserted into destination table 310 or the row provides an update to a corresponding row in destination table 310 (Col. 8, lines 32-36).

The Examiner cites figure 4, step 175, of the Thusoo patent as teaching that input duplicates are appended to the first structure. Applicants respectfully traverse. Step 175 of figure 4 describes updating an existing row. Because an existing row is updated, there is no need to append the input duplicate to the first structure.

The Examiner cites figure 4, step 180 and Col. 8, lines 12-20, as teaching storing index entries for the input duplicates in a second structure. Applicants respectfully traverse. Step 180 of figure 4 describes inserting a row into the destination table (Col. 6, lines 19-20). Col. 8, lines 12-20, describes that a sale amount value from a source table is included in a result set. There is no description of storing index entries in this cited portion of the Thusoo patent.

The Examiner cites figure 5, step 255, and Col. 8, lines 39-41, as teaching automatically re-applying the input duplicates to the first structure. Applicants respectfully traverse. Step 255 describes processing rows that have previously been identified as an update row or an insert row (Col. 7, lines 12-15). Col. 8, lines 39-41, describes that each row in result 340 set may be processed such that the row is either inserted into destination table 310 or the row provides an update to a corresponding row in destination table 310. Merely updating or inserting rows does not teach or suggest automatically reapplying the input duplicates. Furthermore, Applicants describe that re-applying is done by removing the input duplicates from the first structure and applying the input duplicates to matching original rows in the first structure to update the matching original rows, which is not taught or suggested by the Thusoo patent.

The Examiner cites figure 5, step 255, and Col. 8, lines 37-39, as teaching merging index entries from the index entries stored in the second structure to a primary key index. Applicants respectfully traverse. The claim element describes merging index entries in the second structure to a primary key index. The Thusoo patent does not describe storing index entries of input duplicates in a second structure and then later merging these with a primary key index.

The Examiner submits that the Thusoo patent "does not explicitly teach by removing the input duplicates from the first structure and applying the input duplicates to matching original rows in the first structure", but cites the Oracle reference as teaching this. Applicants respectfully traverse. The Oracle reference describes an UPDATE statement, which the Examiner submits is equivalent to removing the input duplicates from the first structure and

applying the input duplicates to matching original rows in the first structure. Applicants respectfully traverse. There is no indication in the Oracle patent that input duplicates are appended to the output table and then re-applied to original rows in the output table.

Claims 7, 17, and 27 describe loading one or more input rows into an output table, wherein input duplicates are appended to the output table, index entries for the input rows are stored in a first structure and discarded input rows are stored in a third structure, wherein the discarded input rows are input rows that are rejected based on a condition (e.g., Specification, page 18, paragraph 60; FIG. 7A, blocks 714-722). Input duplicates are duplicate input rows having a same primary key value, and, of the input rows having the same primary key value, one is an original row. Periodically, the loading of the one or more input rows is interrupted to perform an index merge, wherein input duplicates for which corresponding index entries in the first structure are not added to an index are stored in a second structure (e.g., Specification, page 18, paragraph 61; FIG. 7B, blocks 724-728). It is determined whether to add data for one or more discarded input rows in the third structure to the second structure and, when it is determined that the data for one or more discarded input rows in the third structure are to be added to the second structure, the data for the discarded input rows is added to the second structure (e.g., Specification, page 19, paragraphs 62-63; FIG. 7B, blocks 730-732). Automatically, input duplicates and discarded input rows for which data is stored in the second structure are reapplied to the output table by removing the input duplicates from the output table and applying the input duplicates and the discarded input rows to matching original rows in the output table to update the matching original rows (e.g., Specification, page 20, paragraph 66; FIG. 7D, blocks 742, 744).

The Examiner cites the Abstract and figure 2, Source Table 100, as teaching loading one or more input rows into an output table, wherein input duplicates are appended to the output table, index entries for the input rows are stored in a first structure and discarded input rows are stored in a third structure, wherein the discarded input rows are input rows that are rejected based on a condition. Applicants respectfully traverse. Additionally, Applicants have amended the claims to indicate that input duplicates are duplicate input rows having a same primary key value, and wherein, of the input rows having the same primary key value, one is an original row. The Abstract and figure 2 of the Thusoo patent do not describe appending input duplicates to the output table, storing index entries for the input rows in a first structure, and storing discarded

input rows in a third structure, wherein the discarded input rows are rejected based on a condition.

The Examiner cites figure 4, steps 165-185, as periodically interrupting the loading of the one or more input rows to perform an index merge, wherein input duplicates for which corresponding index entries in the first structure are not added to an index are stored in a second structure. Applicants respectfully traverse. Figure 4, steps 165-185, of the Thusoo patent are directed to processing insert an update rows, and there is no mention of performing an index merge.

The Examiner cites figure 4, step 180, and Col. 8, lines 12-20, as teaching determining whether to add data for one or more discarded input rows in the third structure to the second structure. Applicants respectfully traverse. The claim elements describe a third structure storing discarded input rows, wherein the discarded input rows are input rows that are rejected based on a condition, and a second structure storing index entries. Step 180 of figure 4 describes inserting a row into the destination table (Col. 6, lines 19-20). Col. 8, lines 12-20, describes that a sale amount value from a source table is included in a result set. There is no description of determining whether to add data for one or more discarded input rows in the third structure to the second structure in this cited portion of the Thusoo patent.

The Examiner cites figure 5, step 255 and Col. 8, liens 37-39, as teaching when it is determined that the data for one or more discarded input rows in the third structure are to be added to the second structure, adding the data for the discarded input rows to the second structure. Applicants respectfully traverse. Step 255 of figure 5 describes processing rows that have previously been identified as an update row or an insert row (Col. 7, lines 12-15). Col. 8, lines 37-39, describes that a row with a null value in the key element is preferably inserted into destination table 310. The Thusoo patent describes inserting a row into a destination table, but, this does not teach or suggest the claimed third structure, second structure, and adding the data for the discarded input rows to the second structure.

The Examiner cites figure 5, step 255, and Col. 8, lines 39-41) as teaching automatically reapplying input duplicates and discarded input rows for which data is stored in the second structure to the output table. Applicants respectfully traverse. Step 255 describes processing rows that have previously been identified as an update row or an insert row (Col. 7, lines 12-15). Col. 8, lines 39-41, describes that each row in result 340 set may be processed such that the row

is either inserted into destination table 310 or the row provides an update to a corresponding row in destination table 310. Merely updating or inserting rows does not teach or suggest automatically reapplying the input duplicates and discarded input rows. Furthermore, Applicants describe that re-applying is done by removing the input duplicates from the first structure and applying the input duplicates to matching original rows in the first structure to update the matching original rows, which is not taught or suggested by the Thusoo patent.

The Examiner submits that the Thusoo patent "does not explicitly teach by removing the input duplicates from the first structure and applying the input duplicates to matching original rows in the first structure", but cites the Oracle reference as teaching this. Applicants respectfully traverse. The Oracle reference describes an UPDATE statement, which the Examiner submits is equivalent to removing the input duplicates from the first structure and applying the input duplicates to matching original rows in the first structure. Applicants respectfully traverse. There is no indication in the Oracle patent that input duplicates are appended to the output table and then re-applied to original rows in the output table.

Moreover, the Examiner, in the Response to Arguments, submits that the Thusoo patent in Col. 5, lines 40-44, teaches that the result set may contain each row from Table 1 (source), and that this means that both the source and destination entries will exist in the output table. The Examiner cites this as teaching that input duplicates are appended to the first structure. Applicants respectfully traverse. The Thusoo patent does not address input duplicates, wherein input duplicates are duplicate input rows having a same primary key value, and wherein, of the input rows having the same primary key value, one is an original row.

Applicants respectfully submit that the Thusoo patent or the Oracle reference, either alone or in combination, do not teach or suggest the subject matter of claims 1, 7, 11, 17, 21, and 27.

Dependent claims 2-4, 6, 8, 10, 12-14, 16, 18, 20, 22-24, 26, 28, and 30 incorporate the language of one of independent claims 1, 7, 11, 17, 21, and 27 and add additional novel elements. Therefore, dependent claims 2-4, 6, 8, 10, 12-14, 16, 18, 20, 22-24, 26, 28, and 30 are not taught or suggested by the Thusoo patent or the Oracle reference, either alone or together, for at least the same reasons as were discussed with respect to claims 1, 7, 11, 17, 21, and 27.

Conclusion

For all the above reasons, Applicants submit that the pending claims 1-4, 6-8, 10-14, 16-18, 20-24, 26-28, and 30 are patentable over the art of record. Applicants have not added any claims. Nonetheless, should any additional fees be required, please charge Deposit Account No. 09-0460.

The attorney of record invites the Examiner to contact her at (310) 553-7973 if the Examiner believes such contact would advance the prosecution of the case.

Dated: December 5, 2006

By:_____/Janaki K. Davda/_____
Janaki K. Davda
Registration No. 40,684

Please direct all correspondences to:

David Victor
Konrad Raynes & Victor, LLP
315 South Beverly Drive, Ste. 210
Beverly Hills, CA 90212
Tel: 310-553-7977
Fax: 310-556-7984